Remarks

Applicants appreciate the Examiner's allowance of Claims 19 - 23. Independent method Claims 1, 9, 10, 12 and 14 have been amended to clarify that the claimed methods are for use by a storage switch in a storage network; independent switch Claims 24 and 29 have been amended to clarify that the claimed switches are storage switches for use in a storage network; independent Claims 30 and 31 have been amended to clarify that the claimed linecards are for a storage switch; and independent Claim 34 has been amended to clarify that the claimed software instructions are in a storage switch.

For the reasons which follow it is respectfully submitted that these independent claims, comprising all of the rejected independent claims in the application, and the claims which depend therefrom are allowable over the cited prior art. Accordingly, favorable reconsideration of this application is respectfully requested.

As previously noted, the present invention relates to storage switch architectures, methods and components for use in storage area networks (SANs) that are capable of processing data packets for virtualization without buffering and at wire speed, thereby affording a higher throughput rate than conventional storage switches. As defined and used in the specification, the term "virtualization" refers to the mapping of a virtual target space and virtual addresses to the actual addresses of physical target storage devices. As also defined in the specification (see page 6 paragraph [0015] and page 13, paragraph [0062]), the term "wire speed" refers to introducing no

more latency to a packet than would be introduced by a switch that merely performed switching or routing functions, which means that the switch does not buffer packets. Neither of the cited prior art references discloses or suggests processing of data packets in a storage network to perform virtualization without buffering the packets or at wire speed, as claimed. Accordingly, the cited references can neither anticipate nor render obvious any of the independent claims or the claims that depend therefrom.

The Rejections Under 35 U.S.C. §102(b)

The rejections of Claims 1, 2, 7, 9 and 30 under 35 U.S.C. §102(b) as anticipated by U.S. Pat. No. 6,693,906 to Tzeng are respectfully traversed.

Claims 1, 9 and 30:

Independent Claim 1 has been amended to clarify that the claimed virtualization method is for use in a storage switch in a storage network, and recites in relevant part:

> (a) receiving at a first port of the switch a packet that specifies a virtual target as a destination;
>
> (b) sending at a second port of the switch the packet to a physical target that is associated with the virtual target; and
>
> wherein steps (a) and (b) occur without buffering the packet.

Independent Claim 30 has also been amended to recite that the claimed linecard is for "a storage switch" and the claim recites "means for performing virtualization for a packet without buffering".

Independent Claim 9 is substantially the same as independent Claim 1, except that Claim 9 recites that the claimed method steps occur at "wire speed".

It is respectfully submitted that Tzeng does not disclose or suggest either a virtualization method for use in a storage switch in a storage network, as set forth in independent Claims 1 and 9, or a linecard for a storage switch having means for performing virtualization as set forth in independent Claim 30.

Tzeng does not disclose or relate to a storage switch or to a storage network at all, but rather to a IP network switch for performing layer 2 and layer 3 switching (e.g., Internet protocol processing) in an Ethernet network. As shown in Figure 1 and as described at column 3, lines 45-59 of Tzeng, the switch ports 20 of Tzeng's network switch 12 comprise media access control (MAC) ports that transmit and receive data packets to and from physical layer transceivers at associated network stations. MAC addresses are physical addresses corresponding to actual physical devices and are at the ISO network model physical layer; they are not virtual addresses.

Contrary to the Office's assertion, Tzeng does not disclose or suggest performing virtualization on input data packets to map virtual addresses to physical addresses, or a linecard that performs this function, and certainly does not disclose performing virtualization processing without buffering of packets, as claimed in independent Claims 1 and 30 (and several of the dependent claims).

Tzeng also does not disclose virtualization of data packets at wire speed, as set forth in independent Claim 9 (and several of the dependent claims). Rather, in column 1, lines 39-column 2, line 6, and column 3, lines 60-67, Tzeng discusses the disadvantages of conventional layer 2 and 3 switches that his invention addresses, describes the need to minimizing buffering for higher speed switching, and discloses

that his IP network switch provides layer 2 switching and layer 3 switching capabilities with "minimal buffering". In fact, Tzeng discloses at column 3, lines 60-67 and in Figure 1 that his network switch includes a buffer memory 28 "to store data frames while the switch fabric is processing forwarding decisions for the received data packets". Thus, to the extent that the disclosed switch of Tzeng performs switching based on source address, destination address, and VLAN information, such switching is done using a buffer memory.

Moreover, the "minimal buffering" disclosed by Tzeng is for processing packets for classification according to data format, not for virtualization. (See column 4, lines 27-41.) Tzeng discloses that his Ethernet IP network switch performs packet classification to classify incoming packets by comparing incoming data streams with stored templates to identify data formats relating to certain IP protocols, and uses an associated switch fabric to forward the data packets. In Figure 3, and column 5, lines 35-59, Tzeng discloses that his packet classifier module is configured for "buffer-free identification (i.e., evaluation) of the incoming data packets at the network switch port 20". Thus, the packet classifier module 24 identifies the respective data formats of packets, and while it may do so with minimal or no buffering, the switch does not perform a virtualization function as does the invention, and clearly does not perform virtualization without buffering, as set forth in Claims 1 and 30.

Thus, not only does Tzeng not disclose or suggest a storage switch for a storage network, as claimed, Tzeng further does not disclose or suggest a switch that operates on packets for virtualization without buffering, as recited in Claims 1 and 30.

Moreover, Tzeng also does not disclose or suggest a virtualization method in a storage switch in a storage network that processes data packets "at wire speed" as set forth in Claim 9. At column 8, lines 52-67 referred to by the Office, Tzeng discloses processing incoming data packets using "minimal buffering" so that the network port filter 24 may keep up with the wire rate of incoming data packets reducing the need for buffering. This disclosure has nothing to do with performing virtualization at wire speed, as claimed in independent Claim 9 or dependent Claim 2, for instance, which depends from Claim 1.

Claim 30 Must be Construed Pursuant to 35 U.S.C. §112, ¶6

As to independent Claim 30, which is directed to a linecard for a storage switch, this claim is written in means plus function format and must be construed in accordance with 35 U.S.C. §112, ¶ 6 to cover the corresponding material, structure and acts disclosed in the specification for performing the recited virtualization function, and equivalents thereof. The Office's rejection of Claim 30 fails to properly apply the statutory standard of claim interpretation to Claim 30, and fails to point out either corresponding or equivalent structure in Tzeng to the applicable structure disclosed in the specification that performs the claimed function, or to demonstrate that the Tzeng structure functions as disclosed. Accordingly, the rejection of Claim 30 is improper for this reason also.

In view of the forgoing, it is respectfully submitted that Tzeng cannot anticipate any of Claims 1, 2, 7, 9 and 30, and that the rejections of these claims on Tzeng are improper and should be withdrawn.

The Rejections Under 35 U.S.C. §103

The rejections of Claims 3, 4, 5, 6, 8 and 10 - 35 under 35 U.S.C. §103(a) as unpatentable over Tzeng in view of U.S. Pat. No. 6,400,730 to Latif, et al., are respectfully traversed.

As pointed out above, Tzeng does not relate to a storage switch or network, but rather to an Ethernet switch and data network, and does not disclose or suggest performing virtualization of data packets as claimed, either without buffering of the data packets or at wire speed. As is well known, the problems associated with processing data packets in a storage switch of a storage network are far different and much more stringent than processing of data packets in an IP network. Whereas in an IP network there are procedures in the IP protocol for retransmitting of data packets that are lost or corrupted and reassembly of packet sequence, it is critical in a storage network that storage packets be received in order and without the loss of data packets. Thus, this explains the need for virtualization of storage packets by a storage switch without buffering and at wire speed as done by the invention.

Although Latif discloses a storage network, Latif is concerned with the transferring data over the network between IP network devices and various SCSI or Fibre Channel devices. To handle the diverse devices, Latif teaches converting input IP, SCSI or Fibre Channel protocols to an internal data format handled by his disclosed switch fabric, which performs switching between inputs and outputs, and then reconverting from the internal format back to IP, SCSI or Fibre Channel formats

as appropriate for the device to which the data is being switched. (See column 6, lines 44-57.)

Latif illustrates in Figure 5 a high level switch diagram of his switching device, and describes it beginning at column 7, line 47. As disclosed, switch port interfaces 270 convert data packets between the input frame format and the internal data format for routing through the switch. As described at column 9, lines 48-57 and column 10, lines 26-43, Latif discloses buffers for frame buffering of data to communicate with data frames that fit within an Ethernet link. Accordingly, Latif does not disclose a method, a switch or a linecard that performs virtualization either without buffering or at wire speed, as claimed.

Claims 10 and 12

Independent Claims 10 and 12 are directed to methods for use in a storage switch in a storage network, and are substantially similar. Claim 10 calls for a method for use in a storage switch of a storage network, and recites:

> (a) receiving at a first port located on a first linecard of the switch a packet that specifies a virtual target as a destination;
> (b) the first linecard forwarding the packet to a second linecard of the switch along with information about the virtual target, wherein the second linecard includes a port in communication with a physical target associated with the virtual target;
> (c) the second linecard utilizing the information about the virtual target to update the packet with an address of the physical target;
> (d) sending by the second linecard the packet to the respective physical target; and
> wherein steps (a) - (d) occur without buffering the packet.

Claim 12 is substantially the same except that it calls for a plurality of second linecards which send packets to respective physical targets without buffering the packets.

As discussed above, Tzeng relates to an IP switch and an IP network, not to a storage switch and storage network, and does not disclose virtualization processing of packets without buffering. Furthermore, Tzeng also does not disclose a switch having first and second linecards with respective first and second ports, where the linecards use information about a virtual target to update a packet with an address of a respect physical target, as set forth in Claims 10 and 12. Although Latif discloses a network for transferring data between IP and storage devices, Latif does not disclose linecards that utilize information about a virtual target to update a packet with an address of a respective physical target. Latif also does not disclose processing of packets without buffering. Accordingly, even assuming that Tzeng and Latif may be combined as suggested by the Office (which is not conceded), nothing in the teachings of these two references discloses or suggests the invention of independent Claims 10 and 12, and no logical combination of these references would produce the claimed inventions.

Claim 14

As to independent Claim 14, contrary to the Office's assertion, Tzeng does not disclose receiving a packet at an ingress port (element 20 in Figure 1 of Tzeng) that is destined for a virtual target with a virtual target address. Claim 14 recites, in relevant part, a method for a storage switch comprising:

receiving a packet at an ingress port of an ingress linecard of the switch, said packet destined for a virtual target with a virtual target address;

the ingress linecard retrieving information about the virtual target from a virtual target descriptor, the information including a flowID, and placing a virtual target descriptor identifier and the flowID, in a local header of the packet;

the ingress linecard forwarding the packet to a fabric, which forwards the packet to an egress linecard in accordance with the flowID;

the egress linecard using the virtual target descriptor identifier to identify information about a physical target associated with the virtual target, and using the information about the physical target to convert a virtual target block address to a physical target block address; and

the egress linecard sending the packet to the physical target using the physical target block address.

Ports 20 of Tzeng are MAC ports, as previously pointed out, that receive and transmit data packets to physical addresses of physical layer transceivers. They do not process virtual addresses. Moreover, these ports do not comprise ingress linecards that receive virtual target information including a FlowID, and place a virtual target descriptor identifier in a local header of the packet which is used to forward the packet to a fabric, as claimed. Furthermore, Tzeng does not disclose an egress linecard that uses the virtual target descriptor identifier to convert a virtual target block address to a physical target block address and send the packet to the physical target, as claimed.

Latif also does not disclose either a FlowID or processing of a virtual target descriptor identifier. Figures 6a and 6c of Latif, referred to by the Office, illustrate FCP packet encapsulation required because FCP frames are not directly compatible with an Ethernet interface. The switch of Latif encapsulates FCP packets into an Ethernet frame with a "wrapper" to enable the frame to be processed. (See column 8, lines 10 - column 9, line 1.) This has nothing to do with either a FlowID or a virtual target

descriptor identifier, as claimed, and Latif does not disclose or suggest these elements of Claim 14.

Accordingly, independent Claim 14 and Claims 15-18 which depend therefrom cannot be rendered obvious by the combination of Tzeng and Latif. Moreover, Claims 16 and 17 respectively recite that the process steps are performed at wire speed and without buffering the packet, which also is neither disclosed nor suggested by either of the references for reason discussed above.

Claims 24, 29, 31, 32 and 34

As to independent Claims 24, 29, 31, 32 and 34, Tzeng does not disclose or suggest a storage switch for use in a storage network, a linecard for a storage switch, or a set of software instructions stored in a medium in a storage switch for use by a system for virtualization processing, as claimed.

Tzeng, as previously pointed out, does not disclose or suggest virtualization without buffering. Furthermore, contrary to the Office's assertion Tzeng does not disclose a processor unit in communication with ports to perform virtualization. Rather, as disclosed at column 3, lines 60-64, the CPU 26 of Tzeng controls the overall operations of a switch and the programming of the switch fabric. There is no disclosure or suggestion in Tzeng of the CPU communicating with either a switch port or with another processor unit to perform virtualization, as set forth in Claim 24. Combining Latif with Tzeng also fails to teach or suggest the invention of Claim 24

since Latif fails to disclose or suggest a processor in communication with a port of a linecard for performing virtualization without buffering of a packet, as claimed.

Claim 29 is somewhat similar to Claim 24 but sets forth that the switch comprises a plurality of linecards, each having a plurality of ports and each port being associated with a respective processor unit coupled to a traffic manager that is in turn coupled to a fabric, and wherein the processor unit includes a virtualization unit to translate at wire speed a packet address from a virtual target address to a physical target address.

Independent Claim 31 is somewhat similar to Claim 29 in calling for a linecard for a storage switch wherein a processor associated with a respective port includes a wire speed virtualization unit that includes stored virtual target descriptors and physical target descriptors. Nothing similar is disclosed or suggested by either reference.

Claim 32 is directed to a storage switch, and is drafted in means plus function format. Thus, the claim must be interpreted according to 35 U.S.C. §112, ¶ 6. The Office has not pointed to any structure in either of the references corresponding to the claimed means for receiving at a port a packet destined for a virtual target and sending the packet to a physical target without buffering, and neither reference discloses or suggests any such means.

Finally, Claim 34 is directed to a set of software instructions stored on a medium in a storage switch which are executable by a processor having first and second linecards and corresponding ports which utilize information about virtual

targets to update a packet with a physical target address, and where the execution of the software instructions are without buffering. For the reasons above, the references also fail to teach or suggest the recitations of Claim 34.

Accordingly, it is respectfully submitted that neither Tzeng nor Latif, individually or in combination, teaches or suggests the invention set forth in any of Claims 3, 4, 5, 6, 8 and 10 - 35, and the cited references cannot render these claims obvious.

For the foregoing reasons, it is respectfully submitted that none of Claims 1-18 and 24-35 can be anticipated by or rendered obvious over the combination of Tzeng and Latif, and that all claims are allowable. Accordingly, favorable reconsideration of this application and early allowance of all claims is respectfully requested.

Dated: September 8, 2006

Respectfully Submitted,

_____/Barry N, Young/_____

Barry N. Young
Attorney for Assignee
Reg. No. 27,744

Customer No. 48789
Law Offices of Barry N. Young
Court House Plaza, Suite 410
260 Sheridan Avenue
Palo Alto, CA 94306-2047
Phone: (650) 326-2701
Fax: (650) 326-2799
byoung@young-iplaw.com